

Fast and sensitive read mapping with approximate seeds and multiple backtracking

Enrico Siragusa, David Weese, and Knut Reinert

Department of Mathematics and Computer Science, Freie
Universität Berlin

August 22, 2012

Abstract

We present *Masai*, a read mapper representing the state of the art in terms of speed and sensitivity. Our tool is an order of magnitude faster than RazerS3 and mrFAST, 2–3 times faster and more accurate than Bowtie2 and BWA. The novelties of our read mapper are filtration with approximate seeds and a method for multiple backtracking. Approximate seeds, compared to exact seeds, increase filtration specificity while preserving sensitivity. Multiple backtracking amortizes the cost of searching a large set of seeds by taking advantage of the repetitiveness of next-generation sequencing data. Combined together, these two methods significantly speed up approximate search on genomic datasets. *Masai* is implemented in C++ using the SeqAn library. The source code is distributed under the BSD license and binaries for Linux, Mac OS X and Windows can be freely downloaded from <http://www.seqan.de/projects/masai>.

1 Introduction

Next-generation sequencing (NGS) allows to produce billions of base pairs (bp) within days in the form of reads of length 100 bp and more. It is an invaluable technology for a multitude of applications in biomedicine, e.g. detection of SNPs and large genomic variations, targeted or de-novo genome or transcriptome assembly, isoform prediction and quantification, identification of transcription factor binding sites or methylation patterns. In many of these applications mapping sequenced reads to their potential origin in a reference genome is the first fundamental step preceding downstream analyses.

Because of sequencing errors and genomic variations not all reads occur exactly in a reference genome. Therefore approximate occurrences must be considered and algorithms for approximate string matching tolerating mismatches and indels must be applied to solve the problem. Furthermore, because of homologous and low complexity regions not all reads occur uniquely in a reference

genome. Therefore in some applications, e.g. CNVs calling, all approximate occurrences which could be potential origins must be considered.

1.1 Previous work

All current read mappers can be broadly classified as *best-mappers* or *all-mappers*. Tools in the first class aim at finding the best mapping location for a read according to a scoring scheme eventually taking base quality values into account, while those in the second class aim at enumerating a comprehensive set of locations.

Most prominent best-mappers are based on *backtracking algorithms* for approximate string matching [1]. Substrings of the reference genome within an absolute number of errors from a read are recursively enumerated using a suffix or prefix tree of the reference genome. Since the time complexity of backtracking grows exponentially with the absolute number of errors considered, this method alone is impractical when mapping whole reads with moderate error rates. Hence popular best-mappers [2, 3, 4] apply heuristics to reduce and prioritize enumeration and are optimized to return one or a few best mapping locations.

Conversely, most prominent all-mappers are based on *filtering algorithms* for approximate string matching [1]. Seeds are sampled from given reads and used as anchors to quickly determine, with the help of an index, locations of the reference genome candidate to contain approximate occurrences. Each candidate location is subsequently verified with an online method [5]. Increasing the error rate in filtering algorithms leads to a decrease of the seed length which in turn deteriorates filtration efficiency. Current all-mappers [6, 7, 8, 9] are usually slower than best-mappers but conversely they are able to report all asked mapping locations in reasonable time.

1.2 Our contribution

We present Masai, a read mapper that combines for the first time filtering with backtracking. Our filtering approach is based on non-heuristic and full-sensitive filtration strategies using exact and *approximate seeds*, which are searched in the reference genome via backtracking. Approximate seeds, compared to exact seeds, increase filtration specificity while preserving sensitivity. Moreover, we introduce a *multiple backtracking* method which speeds up filtration by searching all seeds simultaneously with the help of an additional index. Combined together, these methods yield a flexible and efficient filter that significantly speeds up approximate search on genomic datasets.

Masai targets all-mapping, but eventually it can be used as a best-mapper achieving even better runtimes. We extensively compared Masai with popular read mappers on simulated and real datasets. Compared to considered all-mappers, Masai is an order of magnitude faster and has comparable sensitivity. In addition, Masai is more accurate than considered best-mappers and 2–3 times faster than Bowtie 2 and [2] BWA [3]. Masai is implemented in C++ using the SeqAn library and distributed under the BSD license. It can be downloaded from <http://www.seqan.de/projects/masai>.

2 Materials and Methods

In order to map reads to a reference genome, we proceed as follows.

We first construct a conceptual *suffix tree* of the reference genome, then store it on disk and reuse it for each read mapping job. We choose the enhanced suffix array (Esa) [10], which provides an efficient implementation of the suffix tree and consumes 38 Gb of memory for the whole human genome. However, any other data structure equivalent to the suffix tree in terms of allowing a prefix search, i.e. the suffix array [11] or the FM-index [12] can be used to this intent.

At mapping time we choose a filtration strategy according to the reference genome and the specified error rate. Our filtration strategies are based on [13], make use of *exact and approximate non-overlapping seeds* and are guaranteed to be full-sensitive by the pigeonhole principle. In Figure 1 we show an example providing two alternative filtration strategies.

Therefore we partition all reads and their reverse complements in non-overlapping seeds and subsequently arrange all seeds in a conceptual *radix tree*. The time spent to construct the radix tree is easily justified since the tree allows us to perform multiple backtracking. We indeed apply our *multiple backtracking algorithm* to the radix tree, in order to search simultaneously all seeds in the suffix tree of the reference genome.

Finally we perform seed extension on each seed reported by the multiple backtracking algorithm. We extend both ends of each seed using a banded version of *Myers bit-vector algorithm* [14] presented in [6].

In the following of this section we give a detailed explanation of each mapping step.

2.1 Seeds

We now consider formally the read mapping problem. Given a reference genome g , a set of reads \mathcal{R} and an absolute number of errors k consisting of indels and mismatches, for each read $r \in \mathcal{R}$ find all mapping locations where r approximately occurs in g within k errors.

2.1.1 Exact seeds

A simple solution to the problem is provided by a filtering algorithm proposed in [15] which reduces an approximate search into smaller exact searches. Each

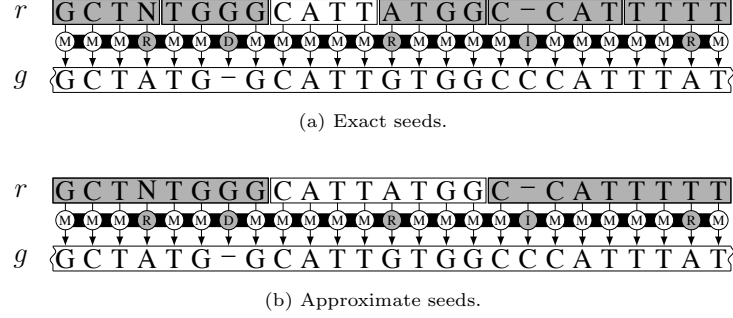


Figure 1: **Filtration strategies.** A read r occurs in the reference genome g within edit distance 5. (a) If we partition r in 6 seeds, at least one seed (in white) occurs exactly in g . (b) Alternatively, if we partition r in 3 seeds, at least one seed (in white) occurs within edit distance 1 in g .

read r is partitioned into $k + 1$ non-overlapping seeds which are searched in g with the help of an index. Since each edit operation can affect at most one seed, for the pigeonhole principle each approximate occurrence of r in g contains an exact occurrence of some seed. However the converse is not true, consequently we must verify whether any candidate location induced by an occurrence of some seed corresponds to an approximate occurrence of r in g .

Filtration specificity in terms of candidate locations to verify is strongly correlated to seed length. Since we want to maximize the length of the shortest seed, we let the minimum seed length be $\lfloor |r|/(k + 1) \rfloor$. If we want to improve filtration specificity by increasing seed length, we resort to approximate seeds.

2.1.2 Approximate seeds

A more involved filtering algorithm proposed in [13] reduces an approximate search into smaller approximate searches. We partition r into $s \leq k + 1$ non-overlapping seeds. According to the pigeonhole principle each approximate occurrence of r in g then contains an approximate occurrence of some seed within distance $\lfloor k/s \rfloor$.

Moreover, we search $(k \bmod s) + 1$ seeds within distance $\lfloor k/s \rfloor$ and the remaining seeds within distance $\lfloor k/s \rfloor - 1$. To prove full-sensitivity it suffices to see that, if none of the seeds occurs within its assigned distance, the total distance must be at least $s \cdot \lfloor k/s \rfloor + (k \bmod s) + 1 = k + 1$. Hence all approximate occurrences of r in g within distance k will be found.

Seeds are searched approximately by backtracking on a suffix tree. We will introduce two efficient multiple backtracking algorithms to search exactly or approximately a set of seeds.

2.1.3 Filtration strategies

With approximate seeds we are free to choose the number of seeds s , which in turn enforces the minimum seed length l to be $\lfloor |r|/s \rfloor$. Or vice versa we fix l , which enforces s to be $\lfloor |r|/l \rfloor$. The resulting filter is flexible, indeed by increasing l filtration becomes more specific at the expense of a higher filtration time.

The optimal seed length l depends on the reference genome as well as on read length and the absolute number of errors. When mapping current NGS datasets on short to medium length genomes, e.g. bacterial genomes, exact seeds are still more efficient than approximate seeds. Conversely on larger genomes, e.g. mammalian genomes, approximate seeds outperform exact seeds by an order of magnitude. Filtration results are provided in the Supplementary Data.

2.2 Indices

We make use of two fundamental data structures, radix and suffix trees. Here we present these indices and give most important implementation details.

2.2.1 Radix tree

The radix tree [16] is a lexicographically ordered tree data structure representing a set of strings. There is one node designated as the root and one leaf per string in the set. Every internal node has more than one child and edges are labeled with non-empty strings. Consequently, common prefixes are compressed and each path from the root to an internal node spells a different substring.

The radix tree for a set of strings can be built in time and space linear in the total length of the strings. It is the ideal data structure to iterate a set of strings in lexicographical order and ask for the longest common prefix of any two strings.

2.2.2 Suffix tree

The suffix tree [17] of a string is the radix tree of all the suffixes of the string itself. It can be built in time and space linear in the length of the string [18].

The suffix tree is used for exact search. A pattern is found by starting in the root node and following the path spelling the pattern. If such path is found, each leaf below the last traversed node points to a distinct occurrence of the pattern in the text.

Approximate search is performed on the suffix tree by means of backtracking [19, 20]. A preorder depth-first search on the suffix tree spells all substrings present in the text. While visiting each branch of the suffix tree, the distance between the pattern and the text spelled along the path is incrementally computed. If the pattern approximately matches the spelled text, each leaf below the last traversed node points to a distinct approximate occurrence of the pattern in the text. Conversely, if the remaining suffix of the pattern can not lead

to any approximate occurrence, the branch is pruned and the visit proceeds on the next branch.

2.2.3 Implementation

We replace the suffix tree with the *enhanced suffix array* (Esa) [10], which preserves the asymptotic performances of the suffix tree and consumes, as implemented in SeqAn, $12n$ bytes for a sequence of length n . We construct the Esa in linear time using the algorithms proposed in [21, 22, 10]. Therefore we use an Esa to index the reference genome. Similarly to all read mappers relying on an index of the reference genome, we build the Esa of the reference genome only once, store it on disk, and reuse it for each mapping job.

We emulate the radix tree by means of the *lazy suffix tree*. We use the *wotd*-algorithm [23] in order to build a partial suffix tree only containing certain suffixes. However, when performing multiple backtracking with exact seeds, the radix tree construction time dominates the overall filtration time. Therefore in this case we resort to the *q-gram index* to emulate the radix tree. We build the *q-gram index* efficiently and in linear time by bucket sort. Below depth q the properties of the radix tree are lost, however multiple backtracking is still applicable.

2.3 Multiple backtracking

We now introduce a method for multiple off-line approximate string matching to search simultaneously a set of patterns in a text. We start by introducing an algorithm for multiple off-line exact string matching and later extend it to approximate string matching.

For simplicity of exposition we describe the algorithms working on tries, although they are easily extendable to work on trees. Hence in the following we assume the text sequence and the set of patterns to be preprocessed respectively using a suffix trie G and a radix trie S . Given a node x , we denote with $label(x)$ the label of the edge entering into x , and with $\mathcal{C}(x)$ and $\mathcal{L}(x)$ respectively the set of children and the set of leaves below x .

2.3.1 Exact search

Algorithm 1 takes as input two nodes g, s respectively of G, S and reports all pairs of leaves $(l_g, l_s) \in \mathcal{L}(g) \times \mathcal{L}(s)$ such that the path from s to l_s spells a prefix of the path from g to l_g .

Consequently by applying Algorithm 1 on the root nodes of G, S we obtain all pairs of leaves (l_g, l_s) such that the pattern pointed by l_s occurs in the text at the position pointed by l_g .

Algorithm 1 Multiple exact search.

```
1: procedure SEARCH( $g, s$ )
2:   if  $s$  is a leaf then
3:     report  $\mathcal{L}(g) \times s$ 
4:   else
5:     for all  $c_s \in \mathcal{C}(s)$  do
6:       if  $\exists c_g \in \mathcal{C}(g) : \text{label}(c_g) = \text{label}(c_s)$  then
7:         SEARCH( $c_g, c_s$ )
8:       end if
9:     end if
10: end procedure
```

2.3.2 Approximate search

Algorithm 2 takes an additional input argument k which denotes the maximum number of mismatches left and computes the union of all paths within k mismatches in the subtrees rooted in g, s . It reports all pairs of leaves $(l_g, l_s) \in \mathcal{L}(r) \times \mathcal{L}(s)$ such that the path from s to l_s spells a prefix of the path from g to l_g with at most k mismatches.

Therefore by applying Algorithm 2 on the root nodes of G, S we obtain all pairs of leaves (l_g, l_s) such that the pattern pointed by l_s occurs within k mismatches in the text at the position pointed by l_g .

Algorithm 2 Multiple approximate search.

```
1: procedure SEARCH( $g, s, k$ )
2:   if  $k = 0$  then
3:     SEARCH( $g, s$ )
4:   else
5:     if  $s$  is a leaf then
6:       report  $\mathcal{L}(g) \times s$ 
7:     else
8:       for all  $c_g \in \mathcal{C}(g)$  do
9:         for all  $c_s \in \mathcal{C}(s)$  do
10:          if  $\text{label}(c_g) = \text{label}(c_s)$  then
11:            SEARCH( $c_g, c_s, k$ )
12:          else
13:            SEARCH( $c_g, c_s, k - 1$ )
14:          end if
15:        end if
16:      end if
17: end procedure
```

For $k = 0$, lines 5–16 of Algorithm 2 are equivalent to Algorithm 1. However Algorithm 1 is preferred to Algorithm 2 because it traverses only edges spelling

common strings instead of all pairs of edges and it is thus more efficient. Figure 2 depicts a run of Algorithm 2.

Algorithm 2 only considers mismatches, but it can be extended to allow indels, e.g. similarly to [13]. In Masai Algorithm 2 is implemented only for mismatches, consequently full-sensitivity is not attained when using approximate seeds and considering mapping locations with indels. However in the results section we show that such implementation detail sacrifices less than 1% sensitivity.

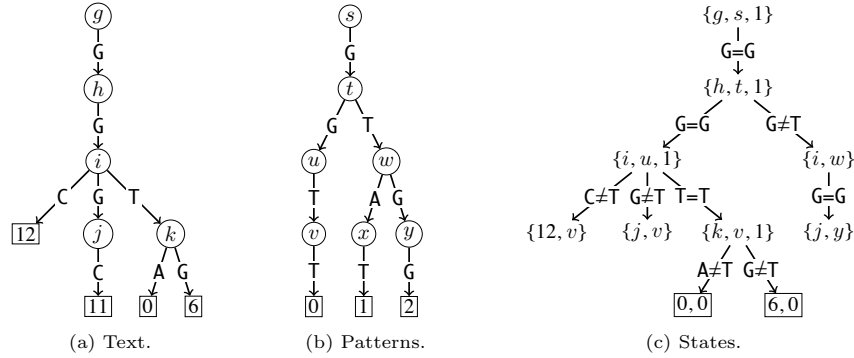


Figure 2: **Multiple backtracking.** (a) A part of the suffix trie representing the text GGTAACGGTGCGGGC (Supplementary Data). Numbers on the leaves are suffix positions in the text, while letters on the inner nodes are arbitrary and serve to distinguish nodes from each other. (b) The trie representing the set of patterns {GGTT, GTAT, GTGG}, respectively numbered {0, 1, 2}. Labels on the leaves show pattern numbers, while labels on the inner nodes are again arbitrary identifiers. (c) Recursive calls performed by Algorithm 2 called with arguments $\{g, s, 1\}$. Edges represent comparisons performed by Algorithm 2 at line 10 or by Algorithm 1 at line 6, nodes with curly brackets represent recursive calls, rectangular leaves represent approximate matches reported. In this example, pattern numbered 0 (GGTT) matches the text twice, at positions 0 and 6, within 1 mismatch.

2.4 Seed extension

We use a banded version of Myers bit-vector algorithm [14] already presented in [6]. Myers' algorithm is an efficient DP alignment algorithm [24] for edit distance. Instead of computing DP cells one after another, it encodes the whole DP column in two bit-vectors and computes the adjacent column in a constant number of 12 logical and 3 arithmetical operations. We implemented a bit-parallel version that computes only a diagonal band of the DP matrix and is faster and more specific than the original algorithm by Myers. More details can be found in the Supplementary Data. However differently from [6], instead

of performing a semi-global alignment to verify a parallelogram surrounding the seed, we perform a global alignment on both ends of a seed. Given a seed occurring with e errors, we first perform seed extension on the left side within an error threshold of $k - e$ errors. Only if the seed extension on the left side succeeds, we perform a seed extension on the right side within the remaining error threshold. Moreover, we first compute the longest common prefix on each side of the seed and let the global alignment algorithm start from the first mismatching positions. We observed that this approach is up to two times faster than [6].

3 Results

We thoroughly compared Masai with the best-mappers Bowtie2, BWA and Soap2 as well as with the all-mappers RazerS3, Hobbes, mrFAST and SHRiMP2. We remark that Bowtie2, BWA, Soap2 and SHRiMP2 rely on scoring schemes taking into account base quality values, while Masai, RazerS3, Hobbes and mrFAST use edit distance. When relevant, read mappers that accept an absolute number of errors (Masai, mrFAST, Hobbes, Soap2) or an error rate (RazerS3) were configured accordingly. We used default parameters, except where stated otherwise (Supplementary Data).

We performed runtime experiments on real data. All read sets are given by their SRA/ENA id. As references we used whole genomes of *E. coli* (NCBI NC_000913.2), *C. elegans* (WormBase WS195), *D. melanogaster* (FlyBase release 5.42), and *H. sapiens* (GRCh37.p2). The mapping times were measured on a cluster of nodes with 72 GB RAM and 2 Intel Xeon X5650 processors running Linux 3.2.0. For running time comparison, we ran the tools using a single thread and used local disks for I/O.

3.1 Rabema benchmark

We first used the Rabema benchmark [25] (v1.1) for a thorough evaluation and comparison of read mapping sensitivity. Similarly to [2], we used the read simulator Mason [26] with default profile settings to simulate from each whole genome 100 k reads of length 100 bp having sequencing errors distributed like in a typical Illumina run. Simulation details are included in Supplementary Data.

The benchmark contains the categories *all*, *all-best*, *any-best*, and *recall*. In the categories *all*, *all-best*, and *any-best* a read mapper had to find all, all of the best, or any of the best edit distance locations for each read. The category *recall* required a read mapper to find the *original* location of each read, which is a measure independent of the used scoring model, e.g. edit distance or quality based. We also classified mapping locations in each category by their edit distance. The benchmark was performed for an error rate of 5%, which corresponds to edit distance 5 for reads of length 100 bp.

Table 1: **Rabema benchmark results.** Rabema scores in percent (average fraction of edit distance locations reported per read). Large numbers show total scores in each Rabema category and small numbers show the category scores separately for reads with $\begin{pmatrix} 0 & 1 & 2 \\ 3 & 4 & 5 \end{pmatrix}$ errors.

	method	all	all-best	any-best	recall
best-mappers	Masai	93.26 <small>99.18 98.73 97.93 95.60 85.77 43.60</small>	97.91 <small>97.79 97.88 98.03 97.98 98.20 96.70</small>	99.95 <small>100.00 100.00 100.00 99.98 99.93 98.71</small>	97.75 <small>97.88 97.84 97.79 97.68 97.56 96.74</small>
	Bowtie 2	92.04 <small>99.18 98.72 96.80 93.44 81.94 40.19</small>	96.16 <small>97.79 97.85 95.80 94.83 93.37 88.86</small>	98.08 <small>100.00 99.96 97.55 96.62 94.93 90.46</small>	95.94 <small>98.01 97.72 95.55 94.24 92.79 89.52</small>
	BWA	92.18 <small>99.18 98.72 97.81 94.25 80.92 37.65</small>	96.81 <small>97.79 97.87 97.88 96.59 92.63 83.47</small>	98.81 <small>100.00 99.95 99.81 98.55 94.28 85.37</small>	96.41 <small>97.93 97.69 97.25 95.77 91.98 84.61</small>
	Soap 2	65.93 <small>99.18 95.55 91.34 8.67 0.70 0.00</small>	69.89 <small>97.79 94.74 91.37 8.96 0.79 0.00</small>	71.37 <small>100.00 96.78 93.18 9.21 0.81 0.00</small>	69.91 <small>98.05 94.62 91.20 11.85 1.41 0.36</small>
all-mappers	Masai	99.90 <small>100.00 100.00 100.00 100.00 99.94 98.58</small>	99.96 <small>100.00 100.00 100.00 100.00 99.93 98.71</small>	99.96 <small>100.00 100.00 100.00 100.00 99.93 98.71</small>	99.96 <small>100.00 100.00 100.00 100.00 99.93 98.71</small>
	Bowtie 2	95.69 <small>99.98 99.91 99.45 97.99 90.69 55.14</small>	98.85 <small>99.74 99.79 98.61 98.21 97.55 93.84</small>	99.16 <small>100.00 99.98 99.01 98.63 97.94 94.17</small>	98.54 <small>99.74 99.58 98.27 97.64 96.87 94.40</small>
	BWA	95.89 <small>99.96 99.88 99.49 97.13 87.79 64.11</small>	97.98 <small>98.81 99.01 99.02 97.83 93.95 85.20</small>	98.82 <small>100.00 99.95 99.82 98.56 94.34 85.37</small>	97.80 <small>99.03 98.96 98.75 97.35 93.43 86.36</small>
	RazerS 3	100.00 <small>100.00 100.00 100.00 100.00 100.00 100.00</small>	100.00 <small>100.00 100.00 100.00 100.00 100.00 100.00</small>	100.00 <small>100.00 100.00 100.00 100.00 100.00 100.00</small>	100.00 <small>100.00 100.00 100.00 100.00 100.00 100.00</small>
	Hobbes	96.56 <small>99.41 99.00 98.76 97.80 93.20 73.05</small>	97.08 <small>97.23 96.59 97.01 97.38 98.16 97.42</small>	98.01 <small>97.92 97.51 97.96 98.43 99.12 98.46</small>	96.41 <small>95.49 95.84 96.54 97.03 97.98 97.79</small>
	mrFAST	99.97 <small>100.00 100.00 100.00 100.00 99.99 99.53</small>	99.97 <small>100.00 100.00 100.00 100.00 100.00 99.10</small>	99.97 <small>100.00 100.00 100.00 100.00 100.00 99.13</small>	99.97 <small>100.00 100.00 100.00 99.99 100.00 99.18</small>
	SHRIMP 2	96.53 <small>99.87 99.82 99.53 98.37 92.58 64.63</small>	99.50 <small>99.34 99.50 99.60 99.64 99.65 98.32</small>	99.85 <small>99.87 99.90 99.91 99.89 99.84 98.57</small>	99.25 <small>99.35 99.30 99.24 99.30 99.09 98.48</small>

For a more fair and thorough comparison we also configured BWA and Bowtie 2 as all-mappers (Soap 2 could not be configured accordingly). To this extent, we parametrized them to be highly sensitive and output all found mapping locations. Since BWA and Bowtie 2 were not designed to be used as all-mappers, they spent much more time than proper all-mappers, i.e. up to 3 hours in a run compared to several minutes. The aim here is to investigate read mapping sensitivity and therefore we do not report running times.

Results for *H. sapiens* are shown in Table 1. Additional results for *E. coli*, *C. elegans* and *D. Melanogaster* are shown in the Supplementary Data.

3.1.1 Best-mappers

Masai showed the best recall values, not loosing more than 3.3% recall on edit distance 5. Conversely, recall values of BWA and Bowtie 2 dropped significantly with increasing edit distance up to loosing respectively 15.4% and 11.5% on edit distance 5. As expected, Soap 2 turned out to be inadequate for mapping reads of length 100 bp at this error rates.

3.1.2 All-mappers

As expected, RazerS3 showed full-sensitivity and mrFAST lost only a minimal percentage of mapping locations. Overall Masai did not loose more than 0.1% of all mapping locations. In particular, Masai was full-sensitive for low-error locations and it lost only a small percentage of high-error locations, i.e. its loss was limited to 0.1% and 1.4% of mapping locations at edit distance 4 and 5. These results show that Masai is suited to replace RazerS3 or mrFAST as an all-mapper.

Table 2: **Variant detection results.** We show the percentages of found origins (recall) and fraction of unique reads mapped to their origin (precision) classed by reads with s SNPs and i indels (s, i) .

		(0,0)		(2,0)		(4,0)		(1,1)		(1,2)		(0,3)	
method		prec.	recl.	prec.	recl.	prec.	recl.	prec.	recl.	prec.	recl.	prec.	recl.
best-mappers	Masai	98.2	98.2	97.6	97.5	96.8	96.8	97.8	97.2	97.9	97.9	97.2	97.2
	Bowtie 2	97.6	97.3	94.6	92.0	92.6	82.5	95.3	93.3	93.5	92.3	96.1	95.4
	BWA	98.2	97.9	97.6	95.3	94.9	85.1	97.4	90.9	97.1	80.3	96.3	66.5
	Soap2	98.1	82.9	97.4	31.0	0.0	0.0	90.6	6.2	0.0	0.0	0.0	0.0
all-mappers	Masai	100.0	100.0	100.0	99.9	100.0	100.0	100.0	99.3	100.0	100.0	100.0	100.0
	RazerS 3	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
	Hobbes	99.9	99.9	99.9	99.9	100.0	100.0	100.0	99.8	100.0	93.6	99.6	90.5
	mrFAST	100.0	99.9	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
	SHRiMP 2	100.0	99.4	100.0	99.7	100.0	99.7	100.0	99.5	100.0	99.2	100.0	99.6

On the other side, BWA and Bowtie2 missed 35% and 45% of all mapping locations at edit distance 5 and their recall values as all-mappers did not substantially increase. Likewise SHRIMP2 could not enumerate all mapping locations, although its recall values were good. Again Hobbes had the worst performance.

3.2 Variant detection

The second experiment analyzes the applicability of Masai and other read mappers in genomic variation pipelines. Similarly to [9], we simulated from the whole human genome 5 million reads of length 100 bp containing sequencing errors, SNPs and indels such that each read had an edit distance of at most 5 to its genomic origin. To distribute sequencing errors according to a typical Illumina run, we used the read simulator Mason. The reads were grouped according to the numbers of contained SNPs and indels, where the class (s, i) consists of all reads with s SNPs and i indels. We mapped the reads with each tool and measured its sensitivity in each class.

We say that a read is mapped *correctly* if a mapping location has been reported within 10 bp of its genomic origin. It is considered to map *uniquely* if only one location was reported by the mapper. For each class we define *recall* to be the fraction of reads which were correctly mapped and *precision* the fraction of uniquely mapped reads that were mapped correctly. Table 2 shows the results for each read mapper and class.

3.2.1 Best-mappers

Among best-mappers, Masai showed the highest precision and recall in all classes. In particular, Masai did not loose more than 3.2% recall in class (4,0), whether Bowtie2 and BWA lost respectively 17.5% and 14.9% and Soap2 was not able to map any read.

Interestingly, we observed that recall values of Bowtie2, BWA and Soap2 were negatively correlated with the amount of genomic variation. For instance, in the Rabema benchmark Bowtie2 lost respectively 7.2 % and 11.5 % of mapping locations at distance 4 and 5, but in this experiment it lost 17.5 % recall in class (4,0). We noticed a similar trend for BWA and Soap2. These tools rely on quality values to guess the best mapping location for a read and tend to prefer alignments which can be explained by sequencing errors instead of true genomic variations. The low performance of Soap2 is also due to its limitation to at most 2 mismatches and no support for indels.

3.2.2 All-mappers

Looking at all-mappers results, Masai showed 100 % precision and recall in all classes, except for classes (2,0) and (1,1) where it lost only 0.1 % and 0.7 % recall. Masai is therefore roughly comparable to the full-sensitive read mappers RazerS3 and mrFAST. SHRiMP 2 showed 100 % precision in all classes but lost between 0.3 % and 0.8 % recall in each class. Hobbes had the lowest performance among all-mappers. It appears to have problems with indels, indeed it lost 9.5 % recall in class (0,3).

3.3 Runtime on real data

In the last experiment we compared the runtime of Masai with the other read mappers. To this end, we mapped the first $10\text{ M} \times 100\text{ bp}$ reads from an Illumina lane of *E. coli*, *C. elegans*, *D. melanogaster* and *H. sapiens*. Whenever possible we asked mappers to map reads within edit distance 5. We measured running times, peak memory consumptions, mapped reads and Rabema any-best scores.

For the evaluation we adopted the commonly used measure of percentage of *mapped reads*, i.e. the fraction of reads for which the read mapper reports a mapping location. However, as some mappers report mapping locations without constraints on the number of errors, we also included Rabema *any-best* scores. The Rabema any-best benchmark assigns a point for a read if the mapper reports at least one mapping location at the minimum edit distance. Final Rabema any-best scores are normalized by the number of reads.

Results for *C. elegans* and *H. sapiens* are shown in Table 3. Additional results for *E. coli* and *D. melanogaster* are shown in the Supplementary Data.

Table 3: **Runtime results.** Results of mapping $10\text{ M} \times 100\text{ bp}$ Illumina reads. **Mapped reads.** In large we show the percentage of mapped reads and in small the cumulative percentage of reads that were mapped with $\begin{pmatrix} 0 & 1\% & 2\% \\ 3\% & 4\% & 5\% \end{pmatrix}$ errors. **Rabema any-best.** In large we show the percentage of reads mapped with the minimal number of errors (up to 5%) and in small the percentage of reads that were mapped with $\begin{pmatrix} 0 & 1\% & 2\% \\ 3\% & 4\% & 5\% \end{pmatrix}$ errors. **Remarks.** SHRiMP 2 was not able to map the H. sapiens dataset within 4 days. Hobbes constantly crashed and was not able to map completely nor the C. Elegans nor the H. sapiens dataset.

dataset		SRR065390 C. elegans				ERR012100 H. sapiens			
		time [min:s]	memory [Mb]	Rabema any-best [%]	mapped reads [%]	time [min:s]	memory [Mb]	Rabema any-best [%]	mapped reads [%]
best-mappers	Masai	3:10	6006	100.00	89.49	24:56	44736	99.99	93.76
	Bowtie 2	24:14	135	99.21	92.58	57:41	3180	99.45	96.72
	BWA	25:53	325	99.33	89.33	80:58	4475	99.54	93.53
	Soap2	4:37	748	95.98	85.95	11:11	5357	95.66	89.73
all-mappers	Masai	10:48	6006	100.00	89.49	284:34	57319	100.00	93.76
	RazerS 3	21:18	11489	100.00	89.49	3653:03	17298	100.00	93.77
	Hobbes	117:46	3885	89.77	80.34	2319:27	71685	59.02	55.35
	mrFAST	67:41	875	99.99	89.49	4462:25	929	99.98	93.75
	SHRiMP 2	541:20	2735	98.51	91.91	—	—	—	—

3.3.1 Best-mappers

On the *C. elegans* dataset Masai was 7.7 times faster than Bowtie2, 8.2 times faster than BWA and 1.5 times faster than Soap2. On the *H. sapiens* dataset Masai was 2.3 times faster than Bowtie2, 3.2 times faster than BWA but 2.2 times slower than Soap2. On one hand, Soap2 was not able to map a consistent fraction of reads because of its limitation to 2 mismatches. On the other hand, Bowtie2 reported more mapped reads than Masai but, taking any-best scores into account, it reported less mapping locations than Masai. On the *C. elegans* and *H. sapiens* datasets, Bowtie2 missed respectively 22.0 % and 20.7 % of reads mappable at edit distance 5. This is due to the fact that Bowtie2 uses a scoring scheme based on quality values and does not impose a maximal error rate threshold.

3.3.2 All-mappers

On the *C. elegans* dataset Masai was 2.0 times faster than RazerS3, 10.9 times faster than Hobbes, 6.3 times faster than mrFAST and 50.1 times faster than SHRiMP 2. Hobbes constantly crashed and mapped less reads than all other mappers in this category. Likewise for Bowtie2, also SHRiMP 2 does not impose a maximal error rate threshold and reported more mapped reads than Masai. However its Rabema any-best score was inferior to Masai. This could be due to the use of a different scoring scheme where two mismatches cost less than opening a gap. Anyway this hypothesis does not explain why SHRiMP 2 did not report some mapping locations at distance 0.

On the *H. sapiens* dataset Masai was 12.8 times faster than RazerS3, 15.7 times faster than mrFAST, and 8.2 times faster than Hobbes which however constantly crashed and mapped only half of the reads. SHRiMP 2 was not able to map the *H. sapiens* dataset within 4 days.

4 Discussion

We showed that, on one hand Masai is faster and more accurate than the best-mappers Bowtie2 and BWA, while on the other hand Masai is slightly slower but substantially more accurate than Soap2. Masai's accuracy becomes considerable in presence of genomic variation, therefore we strongly advise to use Masai in small and large genomic variation pipelines.

At the same time, we showed that Masai is significantly faster than any other all-mapper while being almost full-sensitive. Consequently Masai brings all-mapping within feasible times, although with a higher memory footprint.

In the near future, we plan to index reference genomes using the suffix array or the FM-index to reduce the memory consumption. To achieve full-sensitive mapping on edit distance we will extend multiple backtracking to consider indels.

Masai is implemented in C++ using the SeqAn library. The source code is distributed under the BSD license and binaries for Linux, Mac OS X and Windows can be freely downloaded from <http://www.seqan.de/projects/masai>.

5 Acknowledgements

We thank Manuel Holtgrewe for his joint work on experimental evaluations.

References

- [1] Navarro, G., Baeza-Yates, R. A., Sutinen, E., and Tarhio, J. (2001) *IEEE Data Eng. Bull.* **24(4)**, 19–27.
- [2] Langmead, B. and Salzberg, S. L. (2012) *Nat. Methods* **9(4)**, 357–359.
- [3] Li, H. and Durbin, R. (2009) *Bioinformatics* **25(14)**, 1754–1760.
- [4] Li, R., Yu, C., Li, Y., Lam, T.-W., Yiu, S.-M., Kristiansen, K., and Wang, J. (2009) *Bioinformatics* **25(15)**, 1966–1967.
- [5] Navarro, G. (2001) *ACM Comput. Surv.* **33(1)**, 31–88.
- [6] Weese, D., Holtgrewe, M., and Reinert, K. RazerS3: Faster, fully sensitive read mapping *In press* (2012).
- [7] Ahmadi, A., Behm, A., Honnalli, N., Li, C., Weng, L., and Xie, X. (2012) *Nucleic Acids Res.* **40(6)**, e41.
- [8] Alkan, C., Kidd, J. M., Marques-Bonet, T., Aksay, G., Antonacci, F., Hormozdiari, F., Kitzman, J. O., Baker, C., Malig, M., Mutlu, O., Sahinalp, S. C., Gibbs, R. A., and Eichler, E. E. (2009) *Nat. Genet.* **41(10)**, 1061–1067.
- [9] David, M., Dzamba, M., Lister, D., Ilie, L., and Brudno, M. (2011) *Bioinformatics* **27(7)**, 1011–1012.
- [10] Abouelhoda, M., Kurtz, S., and Ohlebusch, E. (2004) *Journal of Discrete Algorithms* **2(1)**, 53–86.
- [11] Manber, U. and Myers, G. (1990) In SODA : pp. 319–327.
- [12] Ferragina, P. and Manzini, G. (2001) In SODA : pp. 269–278.
- [13] Navarro, G. and Baeza-Yates, R. (2000) *Journal of Discrete Algorithms* **1(1)**, 205–239.
- [14] Myers, G. (1999) *J. ACM* **46(3)**, 395–415.
- [15] Baeza-Yates, R. A. and Navarro, G. (1999) *Algorithmica* **23(2)**, 127–158.
- [16] Morrison, D. R. October 1968 *J. ACM* **15(4)**, 514–534.
- [17] Weiner, P. (1973) In SWAT (FOCS) IEEE : pp. 1–11.
- [18] Ukkonen, E. (1995) *Algorithmica* **14(3)**, 249–260.

- [19] Ukkonen, E. (1993) In CPM : pp. 228–242.
- [20] Baeza-Yates, R. A. and Gonnet, G. H. (1999) In SPIRE/CRIWG IEEE : pp. 16–23.
- [21] Kärkkäinen, J. and Sanders, P. (2003) *ICALP* pp. 943–955.
- [22] Kasai, T., Lee, G., Arimura, H., Arikawa, S., and Park, K. (2001) In CPM : pp. 181–192.
- [23] Giegerich, R., Kurtz, S., and Stoye, J. (2003) *Softw., Pract. Exper.* pp. 1035–1049.
- [24] Needleman, S. B. and Wunsch, C. D. (1970) *J. Mol. Biol.* **48**, 443–453.
- [25] Holtgrewe, M., Emde, A.-K., Weese, D., and Reinert, K. (2011) *BMC Bioinformatics* **12**, 210.
- [26] Holtgrewe, M. Mason – a read simulator for second generation sequencing data Technical Report TR-B-10-06 Institut für Mathematik und Informatik, Freie Universität Berlin (2010).